

IN-60
47326

NASA CASE NO. MSC-20,964-1

PRINT FIG. 3

NOTICE

31P.

The invention disclosed in this document resulted from research in aeronautical and space activities performed under programs of the National Aeronautics and Space Administration. The invention is owned by NASA and is, therefore, available for licensing in accordance with the NASA Patent Licensing Regulation (14 Code of Federal Regulations 1245.2).

To encourage commercial utilization of NASA-owned inventions, it is NASA policy to grant licenses to commercial concerns. Although NASA encourages nonexclusive licensing to promote competition and achieve the widest possible utilization, NASA will consider the granting of a limited exclusive license, pursuant to the NASA Patent Licensing Regulations, when such a license will provide the necessary incentive to the licensee to achieve early practical application of the invention.

Address inquiries and all applications for license for this invention to NASA Patent Counsel, Johnson Space Center, Mail Code AL3, Houston, TX 77058. Approved NASA forms for application for nonexclusive or exclusive license are available from the above address.

{NASA-Case-MS-20964-1) REAL-TIME GARBAGE
COLLECTION FOR IIST PROCESSING Patent
Application (NASA) 31 p CSCL 09B

N87-14863 TSC

Unclas
G3/60 43753

AWARDS ABSTRACT

NASA Case No. MSC-20964-1/
MSC-20965-1

REAL-TIME GARBAGE COLLECTION FOR LIST PROCESSING

The invention relates to a method of practical real time garbage collection for list processing systems.

In a list processing system (FIG. 1), small reference counters CTR are maintained in conjunction with memory cells for the purpose of identifying memory cells that become available for reuse (FIG. 2). The counters are updated as references to the cells are created and destroyed, and when a counter of a cell is decremented to logical zero the cell is immediately returned to a list of free cells. The improvement the invention provides is that when a counter must be incremented beyond the maximum value that can be represented in a small counter, the list element LE represented by the cell A is restructured so that the additional reference count RA can be represented (FIG. 3). The restructuring involves allocation an additional cell AA, distributing counter CTRX, tag TAG TAG2, and pointer CAR CDR information among the two cells, and linking LINK both cells appropriately into the existing list structure B C. Processes for adding references ADD (FIG. 5), deleting references DEL (FIG. 6), and memory retrieval RET and storage STO (FIG. 7) manipulate normal (FIG. 2) and expanded (FIG. 3) cell formats in a manner transparent to the list processor LP.

By the above method, all inaccessible cells are immediately identified and reclaimed; thus there is never an unanticipated delay when needing a free cell. Overhead and complexity are much less than other methods attempting real time garbage collection. Options possible with this method (FIG. 9) provide greater flexibility in cell formats than other methods, which both eases its adaptation to existing list processing systems, and simplifies the design of next generation highly parallel list processing systems. Practical and efficient real time garbage collection as provided by the invention is essential to the anticipated expansion of list processing to support the use of artificial intelligence technology in the areas of monitoring and control, and in the area of reliable conversational interaction.

Inventor:	Robert L. Shuler, Jr.
Employer:	NASA/Johnson Space Center
Initial Evaluator:	Elric N. McHenry

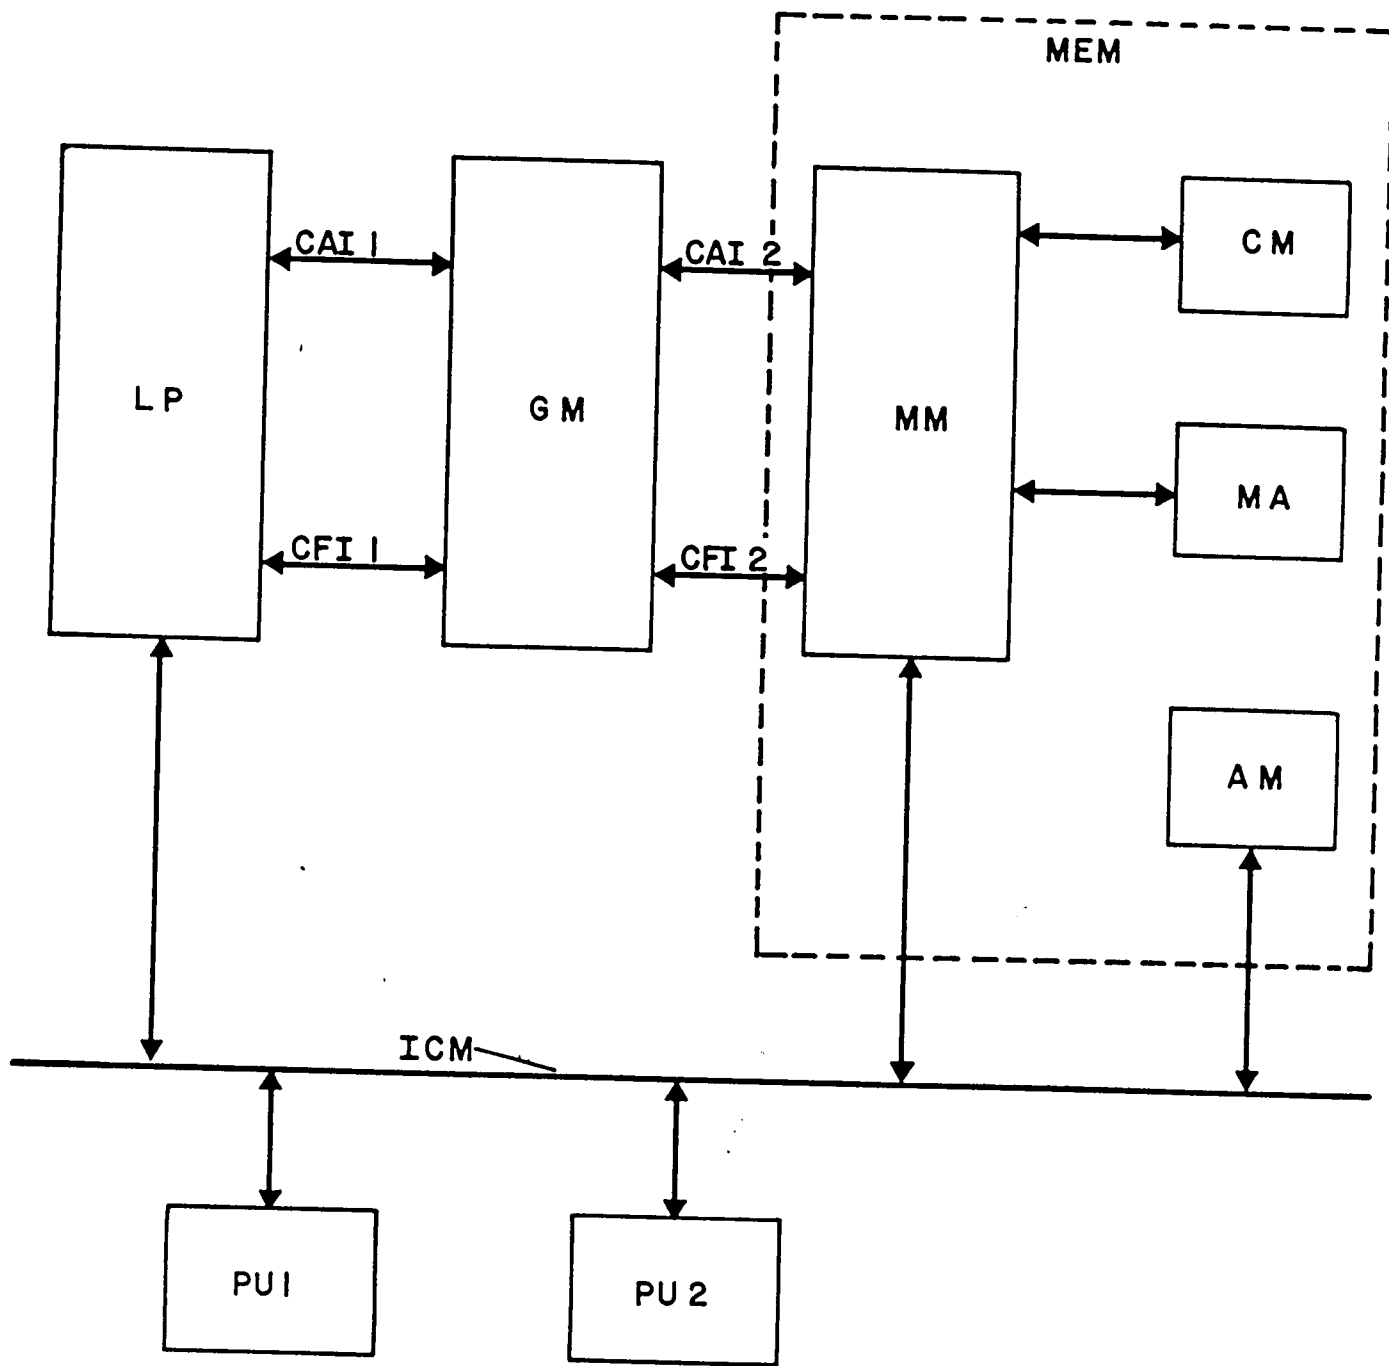


FIG. 1.




FIG. 2.




FIG. 3.

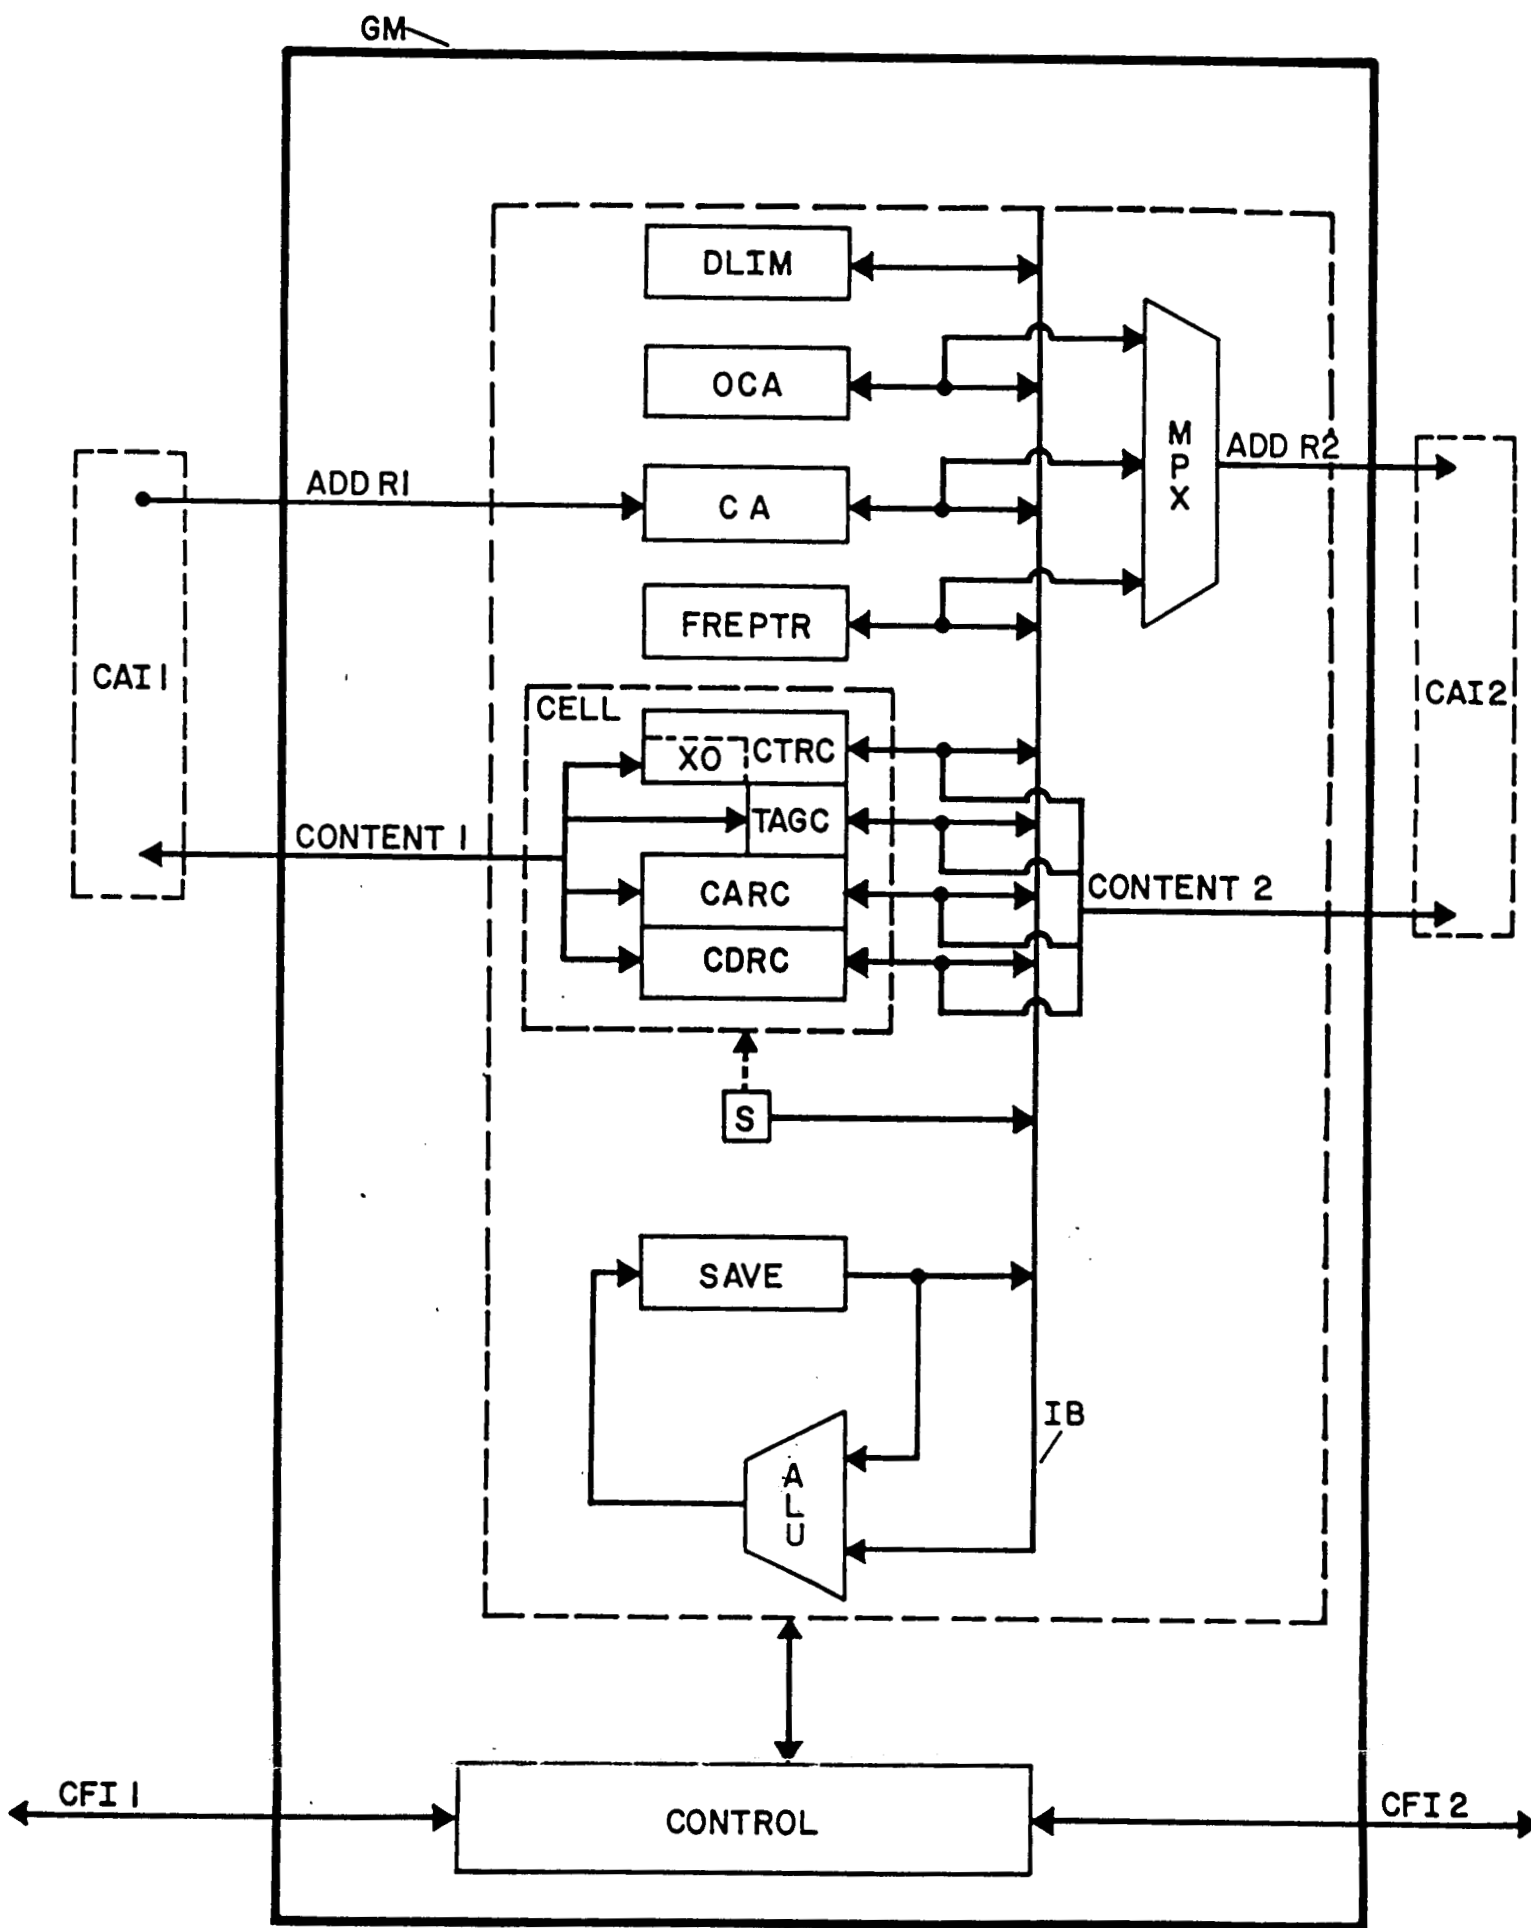


FIG. 4.

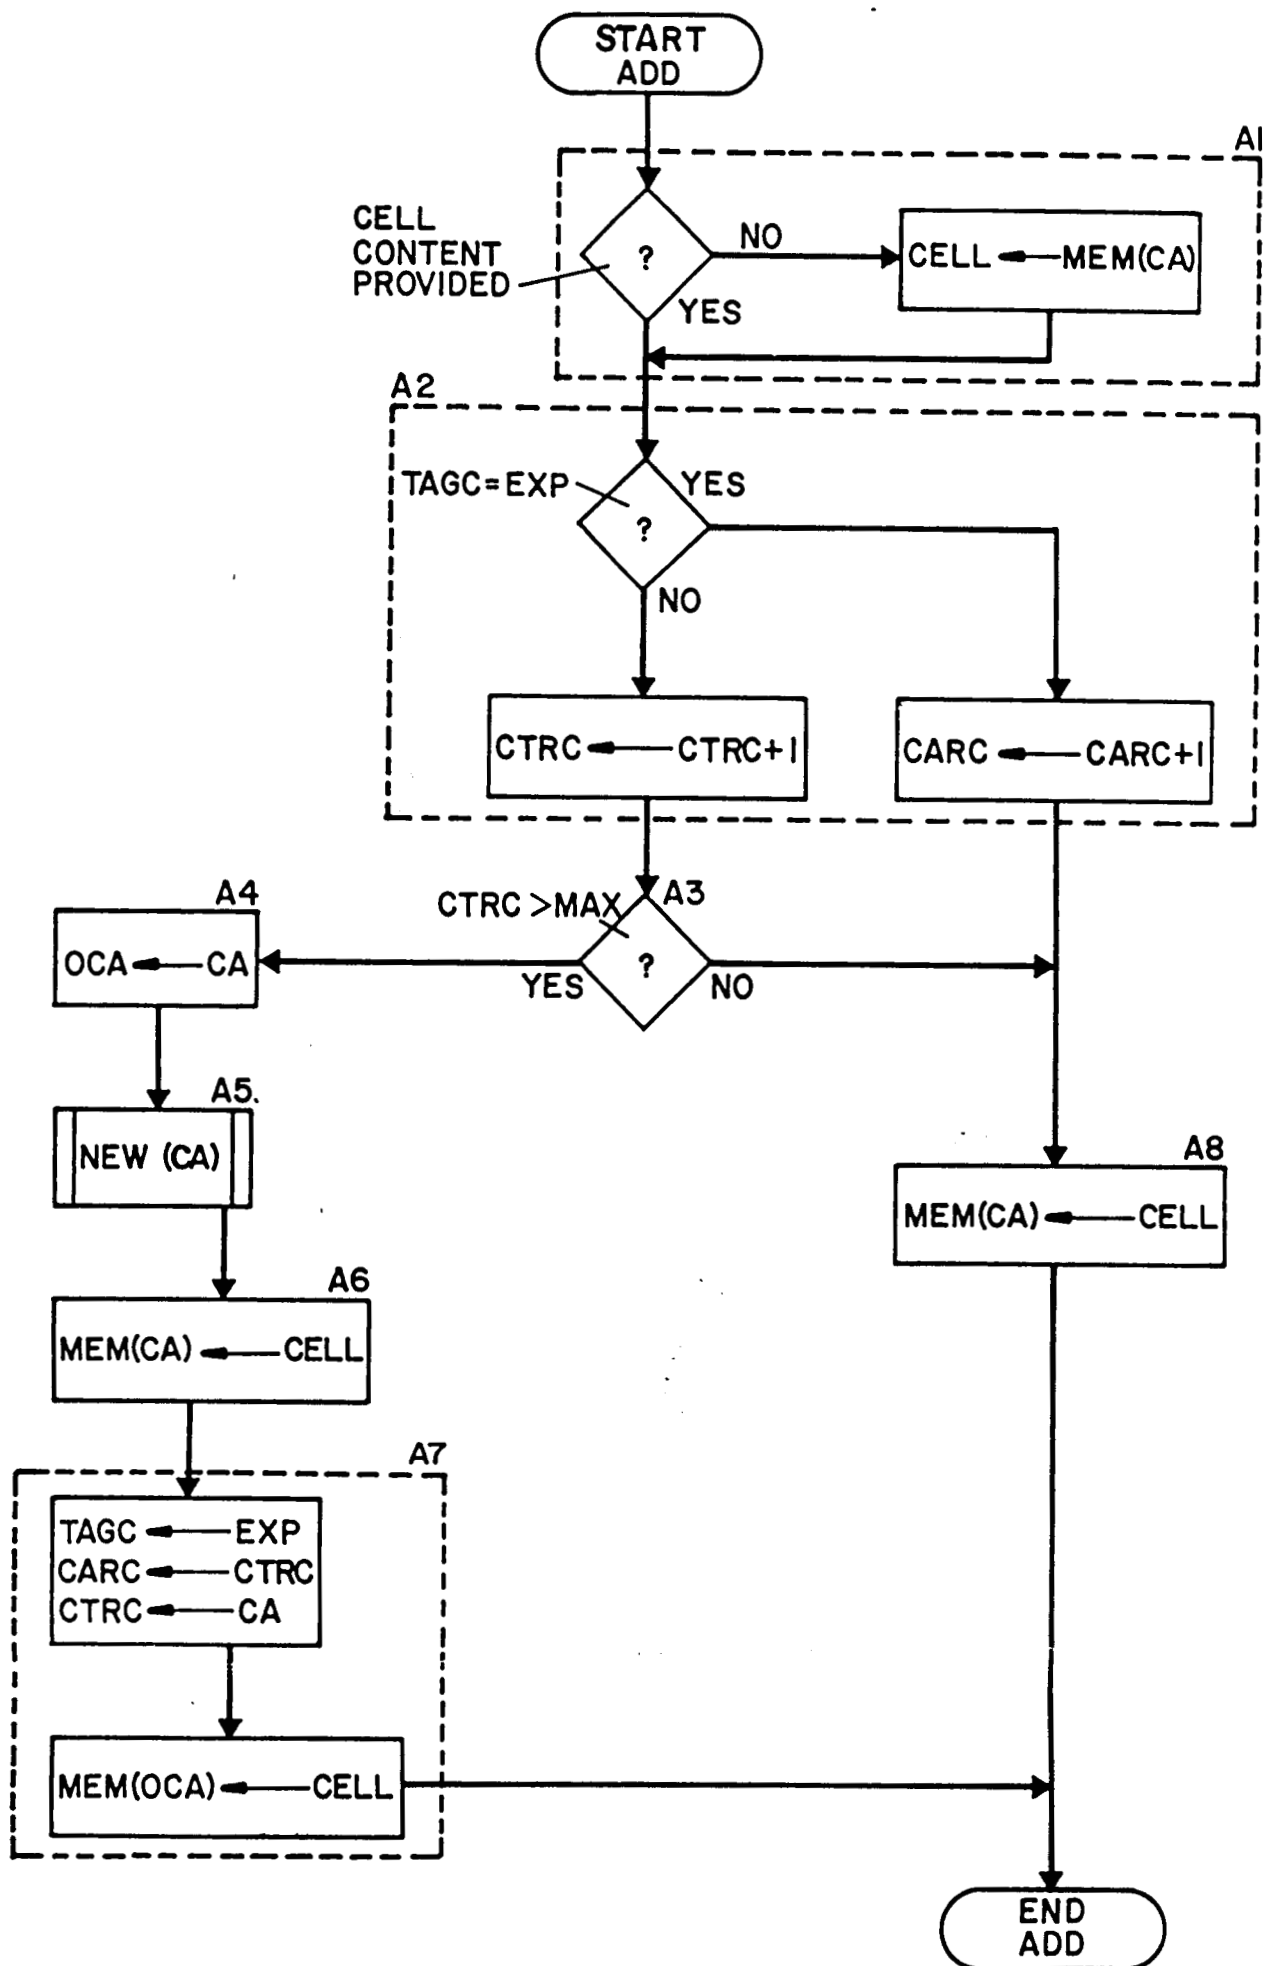


FIG. 5.

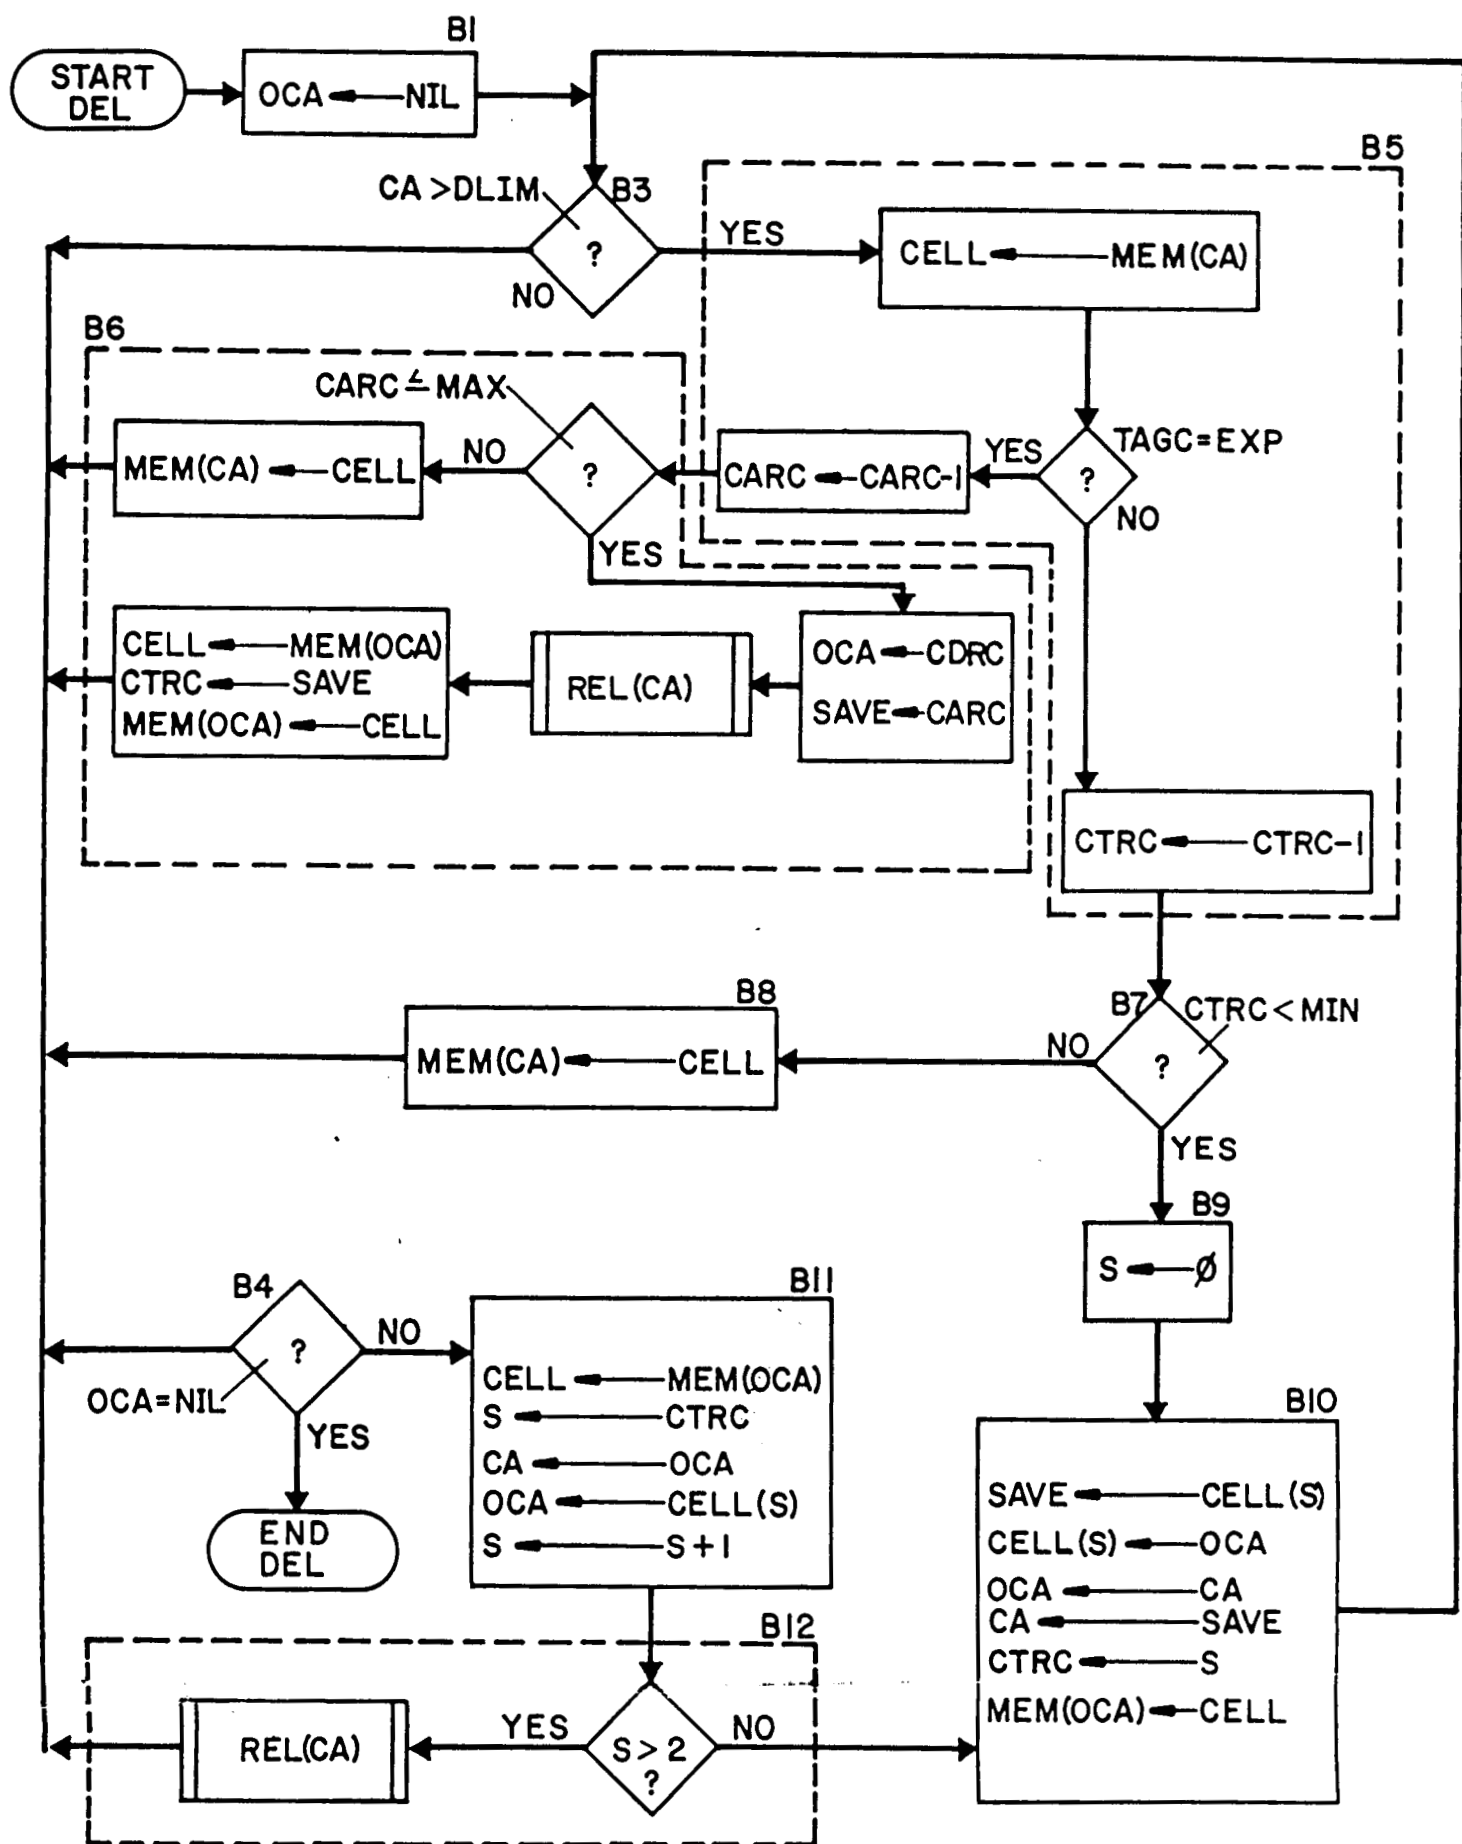


FIG. 6.

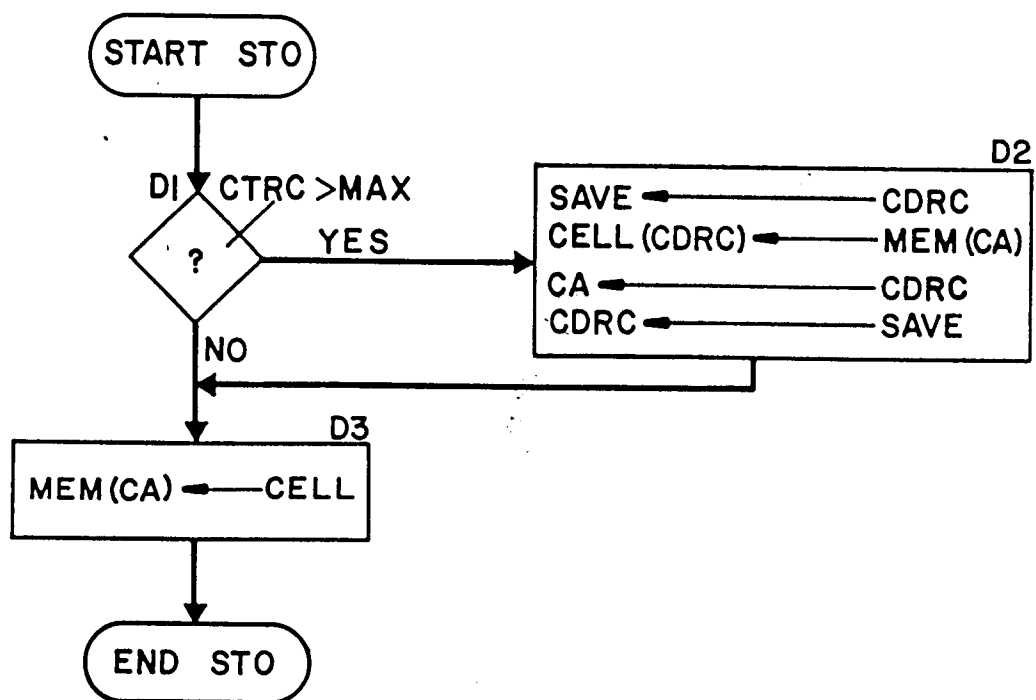
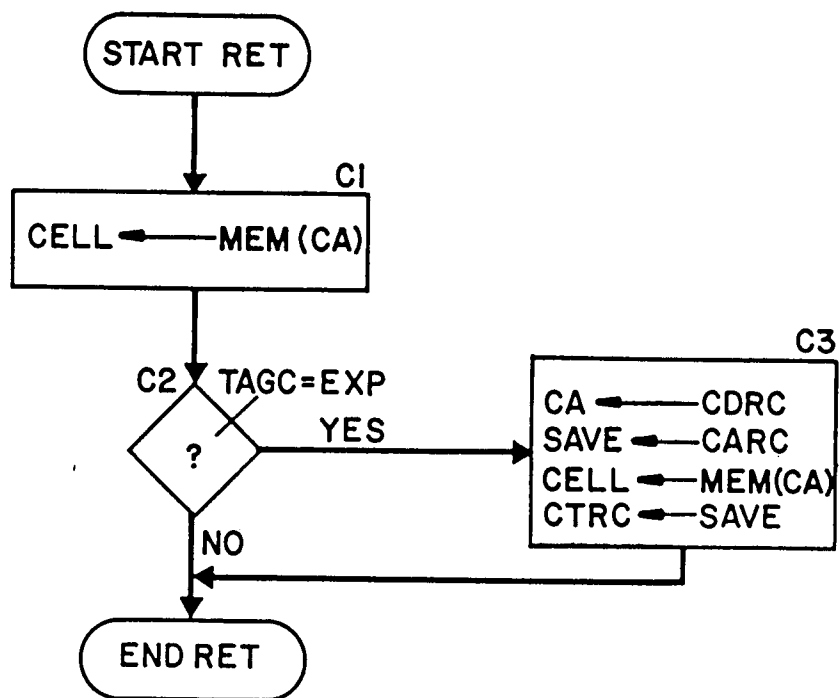


FIG. 7.

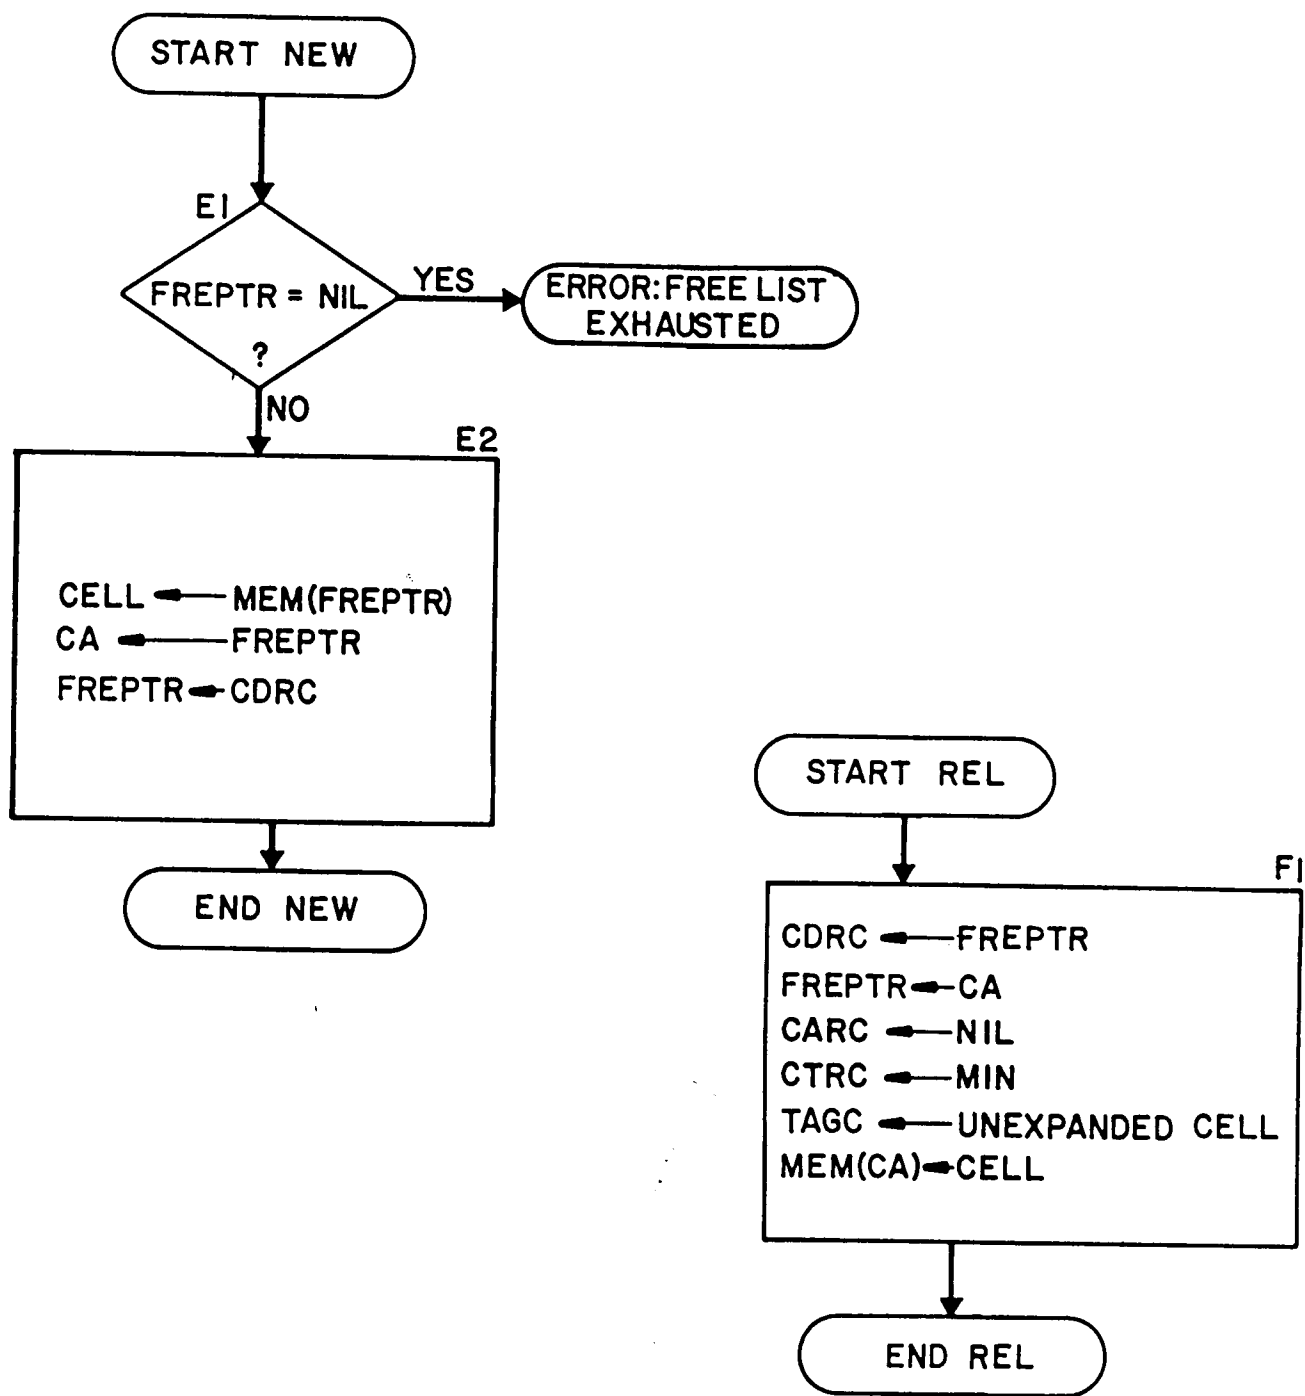


FIG. 8.

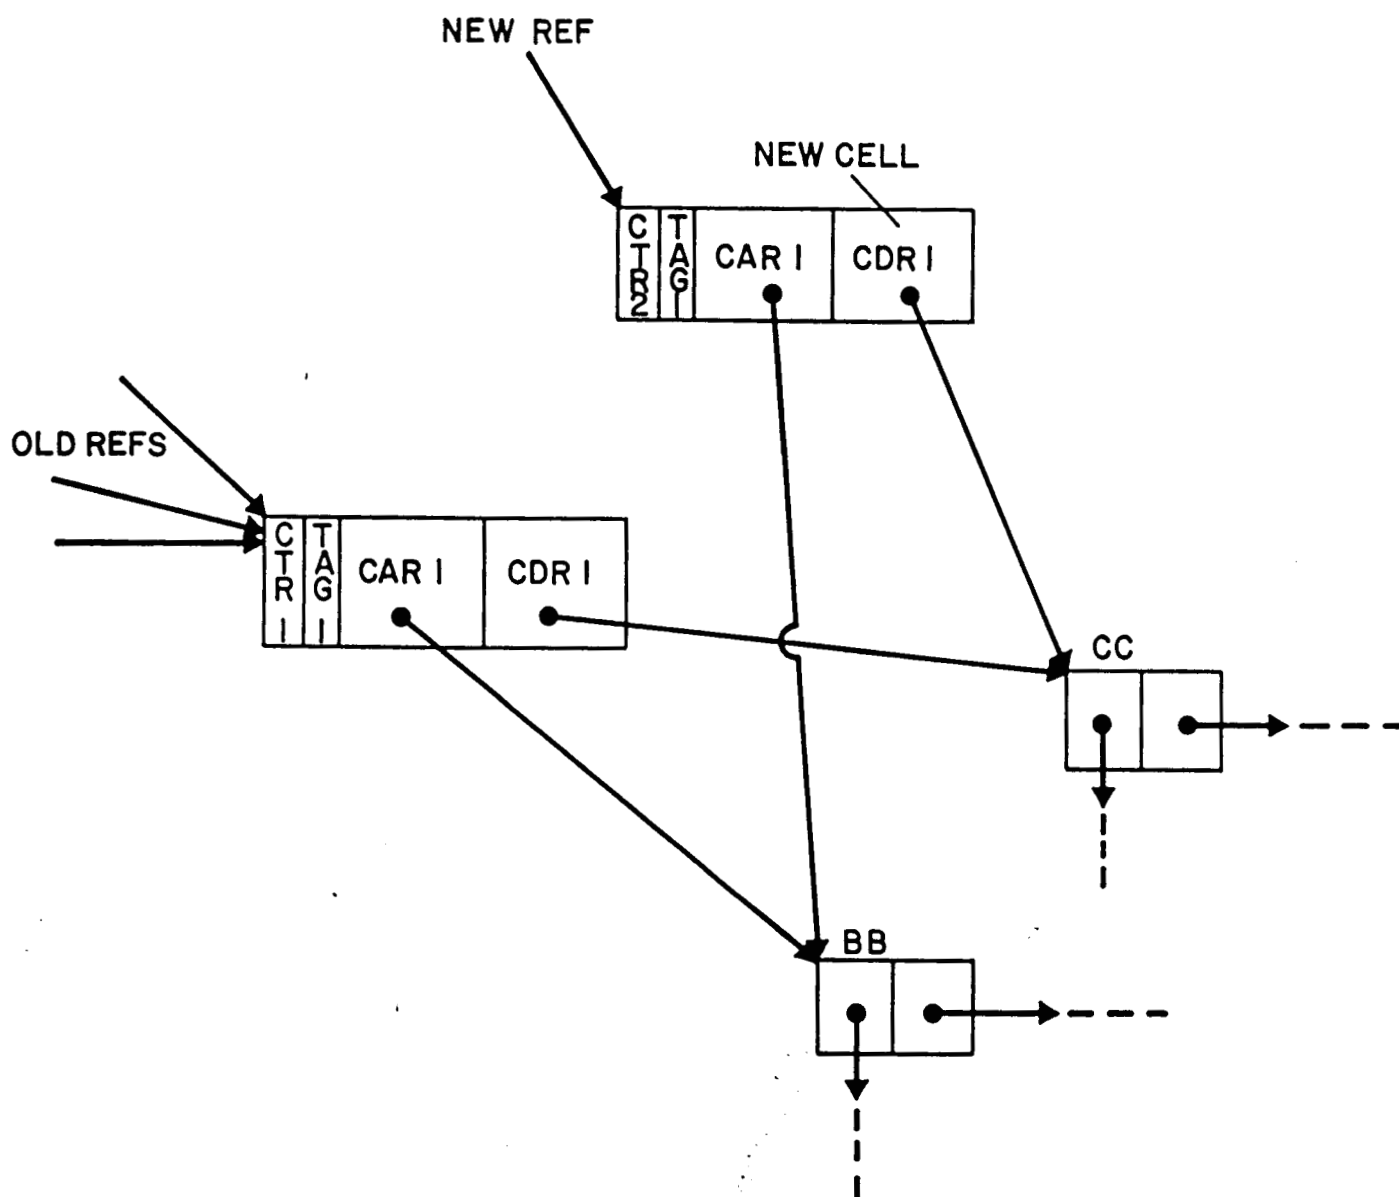


FIG. 9.

-1-

REAL-TIME GARBAGE COLLECTION FOR LIST PROCESSING

Origin of the Invention

The invention described herein was made by an employee of the United States Government and may be manufactured and used by or for the
5 Government of the United States of America for governmental purposes without the payment of any royalties thereon or therefor.

Field of the Invention

This invention relates to data processing systems and their arrangements for allocation and deallocation of memory space, particularly to
10 an improved mechanism for keeping track of the number of active references to a memory cell in a list processing system.

Description of Prior Art

Many present data processing systems are concerned with the manipulation of linked list structures. Each memory cell in a list contains
15 pointers, which refer either to other list fragments, or to fundamental data items which are called atoms. Atoms, which can be symbols or numbers, may also refer to another atom or to a list. New lists are constructed by allocating vacant cells from a free list, and placing
into them pointers to existing lists, pointers to fragments of lists,
20 or pointers to atoms. Pointers within existing lists are not normally modified, and thus several lists or atoms may reliably refer to the same underlying list fragment as part of their value, without having to

-2-

make their own copy. The above described manipulation of linked list structures is termed list processing. It is implemented in specialized data processors designed particularly for list processing, and also in general purpose data processors.

- 5 All accessible memory cells may be reached either by tracing down a list referenced by an atom, by tracing down a list referenced by a stack entry, or by tracing down the free list. As the values of atoms and the stack change, some cells become inaccessible. Identifying these cells and adding them back to the free list is called garbage
10 collection.

In a survey by Cohen, "Garbage Collection of Linked Data Structures," ACM Computing Surveys, September 1981, pp. 341- 367, garbage collection strategies are classified as two main types: (1) mark and sweep, and (2) reference counter based. The basic mark and sweep strategy is to
15 trace down all lists from the base atoms and stack entries, marking each accessible memory cell by setting a bit provided for that purpose. Then memory is scanned, and all unmarked cells are reclaimed. The mark bits are usually also reset during this scan. Processing must be halted while the marking operation is in progress, which can result in
20 large delays. These unanticipated delays cause inconvenience, not to mention outright failure, in systems which must exhibit real time or conversational response, such as process control or spoken natural language communication. In addition to the delay of waiting on the collector to find new free cells, data structures typically become
25 scattered through a large area of memory. In a paging virtual memory system this results in page thrashing, which degrades response time and generally limits the amount of work that can be done by the machine. One improvement to mark and sweep strategies is to use two bits, and a more complicated marking process which is able to proceed without
30 halting the list processor. One such strategy is disclosed in United States Patent #4,121,286 Venton, et. al. However, according to Hickey, "Performance Analysis of On-the-Fly Garbage Collection," Communications of the ACM, Nov. 1984, pp. 1143-1154, up to three times as much processing power may need to be devoted to garbage collection
35 as to list processing in order to guarantee that list processing need never halt to wait for the collector to find a needed free cell.

-3-

A relative of mark and sweep, Baker's Algorithm, is the method used in many commercial list processing systems. This method involves partitioning memory into at least two spaces, evacuating structures from one space to the other, and leaving behind forwarding pointers in the evacuated space. The "to-space" is then purged of all references to the evacuated space via a linear scan in which all pointers to the evacuated "from-space" are replaced with the forwarding pointer. Copying a cell to the "to-space" is equivalent to marking. Another advantage of Baker's algorithm is that cells are allocated sequentially from to-space. A variant of Baker's algorithm is described by Lieberman, "A Real-Time Garbage Collector Based on the Lifetimes of Objects," Communications of the ACM, June 1983, pp. 419-429.

The second method described by Cohen requires keeping a reference counter for each cell, which is incremented when a new pointer to the cell is created, and which is decremented when a pointer is destroyed. When the counter is decremented to zero, the cell may be immediately reclaimed and added back to the free list, thus guaranteeing no delays in finding free cells. Where large cells or blocks of storage are being infrequently manipulated, such as in certain operating system data structures, reference counters have long been used. Their use has not been as common in list processing systems because of the overhead in storing and updating the counters, and because of their inability to reclaim cyclic lists.

Experts disagree over the importance of reclaiming cyclic lists. For example Winston, in his widely used text LISP, 2nd Ed., Addison-Wesley, 1984, p. 141, points out the inadvisability of any structure requiring modification of existing list cells (construction of cyclic lists requires the sort of list modification which renders multiple references to common underlying list fragments unreliable; cyclic structures also render certain processing operations interminable). Lieberman, in the above mentioned article, considers use of cyclic lists to be an important technique.

Overhead is a problem because counters must be theoretically as large as a pointer, and must be kept current. Cohen mentions methods that have been suggested to alleviate one or both the overhead problems for

-4-

reference counters. The earliest is based on the observation that most reference counters will be small; in fact, many will never exceed one or two. In this method, when a counter reaches its maximum value it is no longer updated. When and if memory is finally exhausted, a conventional mark and sweep method is used to reclaim cells with maximum value counters, and to reclaim cyclic lists. United States Patents numbers 4,447,875 and 4,502,118 disclose a very specialized type of list processing system, called a Reduction Processor, having a garbage collection system which uses reference counters in conjunction with mark and sweep.

A more sophisticated method of employing small reference counters, described in Cohen's article, is to assume all cells have a reference count equal to "one," unless the cell is entered in one of several hash tables. The hash table for cells with counts greater than one stores explicitly a counter of necessary maximum size. The tables are not updated immediately, however, due to overhead. Rather, a log of transactions is kept, and the tables are periodically updated; which gets back to the situation of occasional delays. One commercial vendor of list processing machines states that reference counters and tables are used, and these machines exhibit visible pauses for garbage collection.

United States Patent #4,435,766, although not related to list processing or to garbage collection, discloses something which is primitively like a reference counter. This is called a lock counter, and is used to count the number nested resource locks created by a process on a resource, such as a computer peripheral.

Other United States Patents containing teachings of garbage collection in list processing systems, reference counting, replication, cache partitioning, and memory expansion are #4,432,057, Daniell, et. al.; #4,193,115, James Albus; #4,215,397, Gim Hom; #4,558,413, Schmidt and Lampson; and #4,463,424, Mattson and Rodriguez-Rosell.

-5-

Objects of the Invention

It is an object of the present invention to provide an improved reference counter garbage collection mechanism for list processing, which has the advantages of small reference counters, while retaining the
5 absolute determinacy and most of the simplicity of full sized counters.

Additional objects of the invention include: reduction of the overhead of updating reference counters; elimination of memory fragmentation typically caused by mark and sweep methods; and reduction of the complexity and overhead of other reference counter systems attempting
10 to employ small counters.

Another object is to provide these advantages in such a way that they can be incorporated into data processing systems of the type currently in use, with a minimum of impact to the design and operation of these systems.

15 It is also an object of the invention to provide a method of garbage collection which is simple and robust enough to be used in next generation systems, especially those with large memories or employing highly parallel processing.

It is a further object of the invention to provide practical real-time
20 list processing garbage collection.

Further objects and advantages of the present invention will become apparent from a consideration of the drawings and ensuing description thereof.

Summary of the Invention

25 According to the invention, a reference counter of arbitrarily small size is kept for each cell. Each time a new pointer to the cell is created the counter is incremented, and each time a pointer to the cell is destroyed the counter is decremented. When the counter becomes zero the cell is returned to the free list. When any pointers within said

-6-

cell are in turn destroyed, the counters of the cells to which they point are similarly decremented and checked for zero.

On the occasion that a counter can no longer be meaningfully incremented because it has reached its maximum value, an additional cell is
5 obtained. Then the contents of the original cell, some additional count information, and linking information to relate the two cells to the former list structure, are stored in the two cells. The additional count information is incremented to reflect the new reference. The new
reference pointer value will be adjusted to point appropriately within
10 the new cell structure.

By the above method, all inaccessible cells are immediately identified and reclaimed; thus there is never an unanticipated delay when needing a free cell. The fixed and deterministic overhead of updating counters is accepted in lieu of the unpredictable delays of all systems which do
15 not immediately identify and reclaim inaccessible cells. With small reference counters the overhead can be made quite small; less, in fact, than that of mark and sweep systems which either must use a lot of processing power to continuously locate inaccessible cells, or suffer degradation due to memory fragmentation.

20 Description of the Drawings

FIG. 1 is a diagrammatic view of a list processing system showing the invention incorporated therein.

FIGS. 2 and 3 show the structures of a standard cell and an expanded cell, respectively.

25 FIG. 4 is a diagrammatic view of the registers and data paths used by the garbage management system.

FIG. 5 is a flow diagram of the garbage collection algorithm for adding references.

FIG. 6 is a flow diagram of the garbage collection algorithm for
30 deleting references.

FIG. 7 is a flow diagram of the garbage collection algorithm for accessing cells of various types in a uniform manner.

FIG. 8 is a flow diagram of the garbage collection algorithm for obtaining cells from the free list and returning free cells to the free
5 list.

FIG. 9 shows the data structures used to implement an alternate embodiment of the invention in which reference counter information and references to a list structure may be distributed among several memory cells.

10 DESCRIPTION OF THE PREFERRED EMBODIMENT

Referring first to FIG. 1 brief consideration will be given to a typical list processing system organized on a modular basis suited to the invention. The system comprises (i) a central processing unit or list processor LP, (ii) a memory system MEM, (iii) peripheral units
15 PU1, PU2, AM, (iv) a garbage manager GM, and (v) an intercommunication medium ICM for memory to processor or peripheral unit communication. Modules include the provision of needed control information about when references to memory cells are being created and destroyed, and the provision of space within the cell format for storing a reference
20 counter. The arrangement and quantity of the various modules shown in FIG. 1 are typical only and not intended to be limiting.

Interface to the List Processor

The list processor LP is provided with a cell access interface CAI1 for retrieving or updating the contents of memory cells. Such accesses
25 from the list processor LP to the memory system MEM are intercepted by the garbage manager GM, which is interposed between the list processor LP and the memory system MEM. The memory system MEM as shown in FIG. 1 is comprised of a memory manager MM, a cache memory CM, a main memory MA, and an auxiliary memory AM which is typically a peripheral unit
30 such as a disk used as a backing store. Some data processing systems may omit or add elements of the memory system MEM.

-8-

A second interface to the list processor LP is the control function interface CF11 which the processor uses to indicate what type of access to memory is being made, and to perform certain control functions. In addition to a retrieve RTV and a store STR function normally associated
 5 with memory interfaces, there are special control functions which are normally used only by list processors employing reference counter garbage collection. If these special control functions are not already present, the list processor can be appropriately modified to include them in the control function interface CF11. The functions which the
 10 control function interface CF11 communicates to the garbage manager GM are:

RTV - Access to retrieve cell contents

STR - Access to store cell contents

NEW - Get a cell from the free list

15 ADD - Add a new reference to a cell

DEL - Destroy a reference to a cell

EGM - Set the free list pointer and enable garbage manager

DGM - Retrieve free list pointer and disable garbage manager

SDL - Set dynamic space delimiter

20 With each function presented on the control function interface CF11, the list processor LP also provides a cell address on the cell access interface CA11. With access functions, the list processor LP will also provide cell content data (STR), or expect cell content data to be provided to it (RTV). The control function interface CF11 is also used
 25 to return status and exception information to the list processor LP, as for example whether the function was successfully completed, and if not why.

-9-

There is also a cell access interface CAI2 and a control function interface CFI2 from the garbage manager GM to the memory MEM, which are similar to the cell access interface CAI1 and control function interface CFI1, except that the control function interface CFI2 provides
 5 only access (retrieve and store) functions. The cell access interfaces CAI1 and CAI2 and the control function interface CFI2 may be part of the intercommunication medium ICM; however, the control function interface CFI1, because of the various unique functions described above, will be specialized.

10 Division of Responsibility

In a typical list processing system there are several tasks, each with its own logical area of memory. There may also be more than one method of garbage collection available. It is desirable, therefore, that the initiation and termination of the operation of the garbage manager GM
 15 for specified areas of memory be controlled by the list processor LP.

When the list processor LP wishes the garbage manager GM to manage free space in an area of memory, it links the free cells in that area into a free list. If there are pre-existing list structures in the area which were not maintained under garbage manager GM control, the list
 20 processor LP computes and stores correct values for their reference counters. The list processor LP then transmits the address of the head of the free list to the garbage manager GM, along with the enable function EGM, which initiates garbage manager GM control of the free list. After that point, the garbage manager GM assumes all control of
 25 the free list, and the list processor LP retains control of all list elements traceable from atoms and stack entries. The list processor LP may regain full control and retrieve the free list pointer by issuing the disabling command DGM on the control function interface CFI1. The garbage manager GM may also notify the list processor LP of exceptional
 30 conditions, such as free list exhaustion, using the control function interface CFI1.

-10-

Data Structures

FIG. 2 shows the format of a list element LE, comprised of a cell A, to which there are small numbers of references R. A description of each field of bits within cell A is as follows:

- 5 CTR - reference counter having a range of possible values from 1 to the Nth power of 2, where N is the number of bits allocated for the counter.

TAG - a code used by the list processor to indicate the type of cell or other memory data item, in this case an appropriate code to indicate
10 a standard small counter cell.

CAR - the first of the two pointers contained in the cell.

CDR - the second pointer contained in the cell.

FIG 3 shows the same list element as FIG 2, with an additional reference RA, exceeding the capacity of reference counter CTR. As will be
15 explained subsequently, two physical memory cells are now used to represent list element LE. The original cell A has been modified to contain an expanded reference counter CTRX in place of its first pointer CAR, a link pointer LINK to a second cell in place of its second pointer CDR, and an appropriate tag TAG2 to indicate the format
20 of the cell. A second cell AA contains the TAG, CAR, and CDR of the original cell.

Operation of the Garbage Manager

The garbage manager GM is a sequential state machine implementing the process states of FIGS. 5 through 8 as described below. The garbage
25 manager GM has the purposes of maintaining the reference counters and the free list, and of handling memory references on behalf of the list processor LP so that the list processor LP need not normally concern itself with those aspects of cell format which have to do with various reference counter configurations.

-11-

The garbage manager GM has internal storage registers, data paths, and functional units as shown in FIG. 4. When the list processor LP requests a function of the garbage manager GM, it sends the appropriate function code on the control function interface CF11, sends cell
5 address information on the address portion ADDR1 of the cell access interface CA11 to a cell address register CA, and sends and accepts cell content and other information on the content portion CONTENT1 of the cell access interface CA11 to a group of cell content registers CELL, which include: a reference counter CTRC, an extended reference
10 counter portion XC, a tag TAGC, a first pointer CARC, and a second pointer CDRC. Similarly, the garbage manager GM uses the cell address register CA and cell content registers CELL to communicate with the memory manager MM over the control access interface CAI2, along with appropriate function codes on the control function interface CFI2. A
15 memory address can also be supplied from a free pointer register FREPTR, which is used to store the address of the head of the free list, and an old cell address register OCA, which is used in deleting references. A multiplexer MPX is used to select which of these three sources of address information will be sent on the address portion
20 ADDR2 of the cell access interface CAI2. An arithmetic and logic unit ALU is provided for computation and testing. A temporary register SAVE is used for computations and exchanges. Simple transfers are accomplished directly via an internal bus IB. The entire group of cell content registers CELL is transferred on the cell access interfaces as
25 a unit, but one of its component registers is transferred on the internal bus IB. A select register S has the special function of selecting the first pointer CARC or second pointer CDRC for transfer. The delimiter register DLIM is used to partition logical memory space into a dynamic region in which cell allocation is handled by the
30 garbage manager GM, and a static region managed by the list processor LP as will be explained in the discussion of Partial Tag Encoding in Pointers.

FIGS. 5 through 8 define important processes of the garbage manager GM using the functional units of FIG. 4 and the following special terms
35 and conventions:

-12-

EXP - A tag value indicating use of the expanded counter format of FIG. 3.

NIL - A special pointer value designating an empty list.

5 MAX - The maximum reference counter value that can be represented in the small counter format of FIG. 2.

MIN - The minimum counter value representing that only one reference is present.

10 MEM(X)←Y - The operation of storing the contents of a register Y into a cell of memory MEM whose address is in a register X.

Y2←MEM(X2) - The operation of retrieving a cell of memory MEM whose address is in a register X2, and placing the contents of that cell into a register Y2.

15 CELL - Indicates the entire group of registers CTRC, TAGC, CARC, CDRC is referenced or updated, except that when transfer is to or from memory MEM, the extended portion XC of the reference counter CTRC is not included in the transfer.

20 CELL(S) - References the register CARC when the contents of the register S are zero, and references the register CDRC when the contents of S are one.

25 CELL(CDRC) - Indicates transfers which take place as if the entire group of registers CELL were participating, but in which only the register CDRC is allowed to be updated.

NEW(CA) and REL(CA) - Indicate invocation of the obtain cell process NEW and the release cell process REL, which will be described subsequently.

-13-

FIG. 5 defines the process the garbage manager GM uses in response to a request from the list processor LP to add a reference to a cell. The list processor LP must supply a cell address, and may supply the cell contents. Step A1 checks whether cell contents have been supplied, and
5 if not, the garbage manager GM will retrieve them from the memory system MEM. The reference counter is then identified and incremented in step A2. Step A3 checks for small counter overflow. If a previously small format cell's counter becomes larger than can be accommodated within the format, then cell expansion will take place as follows. The
10 garbage manager GM obtains an additional cell from the free list via step A5, which step A6 uses to contain the CAR, CDR, and TAG of the original cell. Step A7 places into the original cell in memory the expanded count, a link to the new cell, and an appropriate tag. Step A8 saves the updated reference counter in memory in the case where cell
15 expansion did not take place.

FIG. 6 defines the process of deleting a reference to a cell. In step B1 the old cell address register OCA is initialized to the value NIL. If in step B3 the cell is found to be not in the dynamic portion of memory, then no further processing of the cell is required, and the
20 terminating step B4 is invoked. At step B4 the old cell address OCA is checked to see whether this deletion was the result of an original request, in which case the process terminates. If in step B3 the reference is to a cell in the dynamic portion of memory (i.e. not an atom), then the cell is retrieved and its counter decremented in step
25 B5. In step B6 the counter portion of an expanded format cell is returned to memory, and expanded counters decrementing below the threshold of expansion cause the cell to be reformatted as a small counter cell, with one of the two cells of the expanded format being returned to the free list. If in step B7 the last remaining reference
30 to the cell has not been deleted then the small format cell is stored in memory via step B8, otherwise the cell must be returned to the free list. Returning the cell to the free list requires deleting any references which the cell makes to other cells, a process handled entirely within the garbage manager GM. This recursive function is
35 accomplished without a stack by using the cells being freed to store information which is local to each level of recursion. The CA register contains the address of the cell of current interest. If there was a

-14-

previous cell, its address is in OCA. A still prior cell address is stored in the cell addressed by OCA. The S register is used to indicate which pointer within the current cell is being processed, the CAR or CDR. When a cell is to be freed, then S is set to zero in step B9, 5 which selects the CAR. In step B10 an exchange is then performed in which the old cell address OCA is moved into CELL(S), the current cell address CA is moved to OCA, the former contents of CELL(S) are moved to CA, which will become the new cell address of interest, and the value of S itself is saved in the counter field of the current cell. The 10 current cell is then stored back to memory so that the S and OCA values in it, as well as the CDR pointer, may be recalled when needed. The process of considering the current cell address in register CA as a deleted reference then begins again with step B3. When such process is finished, the value in register OCA is used to determine whether it was 15 an initial deletion requested by the list processor LP which has finished, or whether it is a deletion that was invoked by the garbage manager GM. In the latter case, the OCA register is used in step B11 to retrieve the former cell of interest, whose contents are used to restore other necessary information that was saved earlier. Then S is 20 incremented, and it selects the CDR of the current cell for deletion. When control is again returned to step B11, incrementing S reveals neither CAR or CDR to be selected, so the current cell is ready to be returned to the free list via step B12, and its handling is complete.

FIG. 7 defines how the garbage manager GM responds to requests from the 25 list processor LP for cell storage and retrieval. On a retrieval function RET the cell contents are obtained from memory in step C1. If in step C2 the cell turns out to be in expanded format, then the second cell of the pair is also be retrieved, and the information it contains is passed back to the list processor LP. On a store function STO step 30 D1 determines whether or not the cell is in expanded format by looking at the count value of the cell, which is always maintained to full precision in communications between the garbage manager GM and the list processor LP. If the cell is in expanded format, then the first member of the cell pair is retrieved in step D2 in order to obtain the address 35 of the second cell of the pair, which is then used by step D3 to store the CAR, CDR, and TAG from the list processor.

-15-

FIG. 8 defines the processes of obtaining a cell from the free list, NEW, and of releasing a cell to the free list, REL. These processes may be invoked by the list processor LP by using the function codes for retrieval RTV and storage STR on the control function interface CF11, or by other garbage manager GM processes. In the obtain cell process NEW, step E1 checks for possible free list exhaustion, and step E2 obtains the address of the first cell from the free list, putting that address in the cell address register CA for communication back to the invoking process. In the release process REL, the cell to be released is threaded on to the head of the free list by updating its pointers and updating the free list pointers as shown in step F1.

Cache Operation

While correct logical function of the garbage manager GM is not dependent on any particular implementation of the memory subsystem, its efficiency is. As seen from the preceding process descriptions, the garbage manager generates additional memory references, many of which are store operations. References to the same cell are frequently close together in time. Therefore, if the memory subsystem uses a high speed cache buffer having the characteristic that every update operation is not written to main memory (i.e. main memory is updated only when the contents of that particular cache cell must be evacuated to hold another memory cell), then overall performance will be greatly improved.

Addition and Deletion of References by the List Processor

The list processor LP exercises a great deal of control over the efficiency of the garbage manager GM by the frequency with which it requests addition and deletion of references. Whenever the list processor performs a modular operation over a list structure which is static for the duration of the operation, however complex that operation may be, then reference control requests may be deferred until the end of the operation. This results in the elimination of many intermediate reference control operations. For example, consider a list processing primitive which scans a list looking for a particular item. Each operation in updating a list scanning pointer to the next element in

-16-

the list could be viewed as requiring one reference deletion and one reference addition. Alternately, knowing the structure of the operation being performed, it becomes necessary to perform only one reference addition (for the result at the end of the operation), and one
5 deletion (for the initial argument structure, again performed at the end of the operation). To go even further, reference addition and deletion in the above example can be made the responsibility of whatever routine invoked this function, allowing that routine to also optimize its reference control operations.

10 Partial Tag Encoding in Pointers

A further efficiency consideration concerns the ability to determine whether a referenced memory item is a dynamically allocated cell, or a static entity such as an atom, by examining the pointer to the item. This may be done, for example, by partitioning the address space into
15 static and dynamic portions as described above, which is particularly convenient in virtual memory or segmented memory systems. If such is the case, then addition and deletion of references to static items will not require additional memory references. If such is not the case, then the items will have to be retrieved and their tag fields examined even
20 if they are static. Stack entries are considered static for this purpose.

DESCRIPTION OF ALTERNATE EMBODIMENTS

The embodiment described above has the advantage that it easily interfaces with certain types of existing list processing systems. Those
25 skilled in the art will recognize various alternate embodiments, some of which are more suitable for their purposes. Selected ones are briefly described below.

Software Implementations

Dynamic expansion of reference counters could be emulated by list
30 processing software running on a conventional data processor. This has been accomplished to verify the concepts and principles of the method of garbage collection set forth above. Software implementation also

-17-

has a use in studying the behavior of the garbage manager for alternate configurations of reference counter size and expansion format.

Zero Size Reference Counters

- The small counter format cell may be so structured that it has no space allocated for a reference counter in which case it is presumed to have the value one. When a reference to such a cell is deleted, then the cell is returned to the free list. When a reference is added, the cell must be expanded. The effectiveness of such small counters depends upon the observation that a majority of counters have the value one in many list processing applications, and upon the ability of a cache memory to handle temporary excursions above the value one without actually expanding and contracting the cell in memory. The cache might, for example, employ a third counter size chosen to handle most such excursions.
- Using such a counter size, the reference counter method could be more easily adapted to a list processing system which does not have any bits reserved for garbage collection purposes. This includes some systems which use Baker's algorithm. A second principle advantage of a zero count system is that it allows all non-data bits, such as tags, to be moved out of the cell and into the reference (pointer) to the cell. In mark and sweep garbage collection this cannot be done because the cells are accessed during the sweep phase by a scan of memory independent of the pointers to the cell. In a normal reference counter system it cannot be done because the counter itself must be present. To remove all such non-data bits, fully encoding the tag in the pointers to the cell, has the advantage that the type of cell is known from the pointer without having to retrieve the cell, and the advantage that cell data content may use the full memory cell size. Standard 32 bit data formats could, for example, be used in a processor employing a common 32 bit memory width.

Other Arrangements of CTR, TAG, CAR, CDR and LINK

When a counter must be expanded, there are many ways of allocating the cell information among the two cells. In addition to just placing the

-18-

information differently than in FIG. 3, the counter information may be distributed between the two cells. FIG. 9 shows an expansion in which an original cell OLDCELL is left completely unmodified by the expansion, and a new cell NEWCELL contains a new counter CTR2 of the same size as the old counter CTR1. The added reference NEWREF is adjusted to point to the new cell NEWCELL. The pointers CAR1 and CDR1 of the new cell NEWCELL are copied from the original cell OLDCELL. In this way the link is from the new cell NEWCELL to the list structure BB and CC being referenced by the original cell OLDCELL, rather than between NEWCELL and OLDCELL. This distribution has the result that no reference is added to the original cell OLDCELL, and its reference counter CTR1 may remain at the same value. The new reference NEWREF is to the new cell NEWCELL. The new cell NEWCELL then adds new references to two other already existing cells BB and CC, whose reference counters must be incremented, and which may of course have to be expanded if their reference counters are already at maximum value. In the worst case the entire structure being referenced has all its reference counters at maximum value, and thus the entire structure is copied through individual expansions of each of its cells.

20 This distribution of counter information among several small counters has the advantage of maintaining a uniform cell format, and of eliminating the extra retrieve operations to get the second member of an expanded cell pair. Its disadvantage is that list processing software which employs list splicing techniques would need to be carefully examined to assure that it would produce the anticipated result.

Strategies may be mixed. For example, zero size counters may be maintained for dynamically allocated numeric quantities resulting from computation, while small counters of some other size are used for list cells containing pointer pairs. Any of the distribution schemes, or a mix in which some cells are expanded one way and some another, may be used with the various cell types.

Addition to a Conventional Data Processor

The function of the garbage manager GM may be placed on the memory bus of a conventional data processor, in a manner similar to a memory

-19-

module or peripheral controller. It may include its own memory, or re-direct references again on the bus to the system's memory. Since there are no dedicated signal paths for the reference control information, it would be communicated by some other means, as for example by
5 storing a special code in a fixed address, or by accessing one of several fixed addresses. The "store immediate" instructions of some processors would be suitable for this purpose. Such an embodiment would allow efficient use of the invention in conjunction with a conventional processor.

10 Closely Integrated Processor and Garbage Manager

The list processor and garbage manager may share data paths, functional units, and sequencers. This would require a close coupling of the two, but could produce an economic embodiment for purposes such as implementation of a list processor on a VLSI (Very Large Scale Integration)
15 chip.

Multiple Processors and Highly Parallel Processors

Where there are multiple processors and each has its own memory, each would also have its own garbage manager. The simplicity and determinacy of garbage management using the present invention would permit
20 simpler processors and would make coordination among the processors easier. The immediate identification and reuse of garbage cells minimizes the amount of memory required for each processor.

Where there are memory modules separate from the processors, with some means of interconnecting the processors and the memories, a garbage
25 manager could be included either with each processor, or with each memory module. In the case of including a garbage manager with each processor, some means would need to be provided to assure consistent results when two or more processors were updating elements of the same memory module. In the case of including a garbage manager with each
30 memory module, interconnection traffic would be reduced (because expansions and second cell accesses are handled locally), and the problem of synchronizing multiple access would be somewhat reduced.

-20-

The above configurations avoid the problem typically encountered of having to scan the pointers of all other memory modules when looking for garbage within a particular module. This becomes more important as memories become larger and are partitioned into more modules to support
5 parallel processing.

Garbage management in the manner prescribed by the invention is also compatible with methods of controlling the sharing of transient list structures, such as copying lists, or use of a forwarding table. The garbage manager may even be used to implement the operation of a
10 forwarding table by merely marking the table entries as being in expanded format, and by providing some means to inhibit the de-expansion of table entries (which could be as simple as initializing each entry with a count exceeding the maximum small format counter value).

15 Those skilled in the art will recognize that many other embodiments may be found which use the basic principles of the invention.

21
~~25~~

ABSTRACT

In a list processing system, small reference counters are maintained in conjunction with memory cells for the purpose of identifying memory cells that become available for re-use. The counters are updated as
5 references to the cells are created and destroyed, and when a counter of a cell is decremented to logical zero the cell is immediately returned to a list of free cells. In those cases where a counter must be incremented beyond the maximum value that can be represented in a small counter, the cell is restructured so that the additional refer-
10 ence count can be represented. The restructuring involves allocating an additional cell, distributing counter, tag, and pointer information among the two cells, and linking both cells appropriately into the existing list structure.